

Europäisches Patentamt
European Patent Office
Office européen des brevets



(11) EP 0 802 485 A1

(12) EUROPEAN PATENT APPLICATION

(43) Date of publication:
22.10.1997 Bulletin 1997/43

(51) Int. Cl.⁶: G06F 12/08, G06F 11/10

(21) Application number: 97104950.7

(22) Date of filing: 24.03.1997

(84) Designated Contracting States:
DE FR GB IT SE

(30) Priority: 15.04.1996 US 631934

(71) Applicant: SUN MICROSYSTEMS INC.
Mountain View, California 94043-1100 (US)

(72) Inventors:
• Gittins, Robert S.
Woodland Park, CO 80866 (US)
• Corbin, John R.
El Paso, TX 79912 (US)

(74) Representative: Hanna, Peter William Derek et al
Tomkins & Co.,
5 Dartmouth Road
Dublin 6 (IE)

(54) Host-based RAID-5 and NV-RAM integration

(57) In a computing system utilizing redundant storage devices arranged in the RAID disk array (44), data is stored in the computing system using a memory cache (42) created from system memory and the disk array. A checkpoint module (90) detects a fault in the computing system and generates a fault indication, and a cache manager (92) writes data and parity to the

memory cache (42) in a first mode, and writes data and parity to the storage device (44) in a second mode. In response to the fault indication (100, 102, 104), the checkpoint module copies the data contained in the cache to the disk array, and switches the cache manager from the first mode to the second mode.

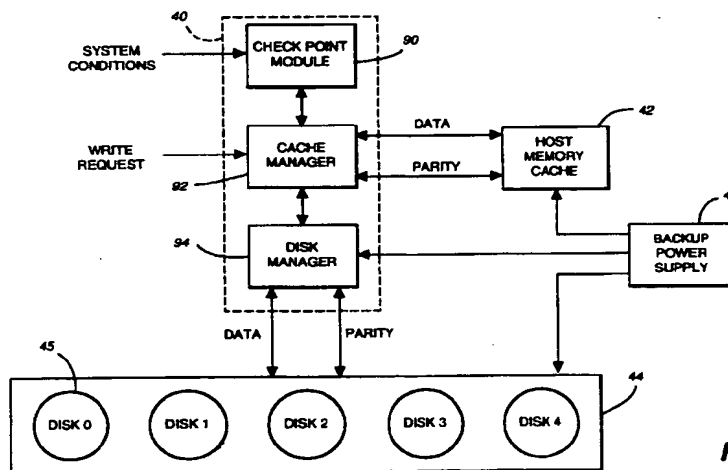


Fig. 5

EP 0 802 485 A1

Description

Field of the Invention

The present invention relates, in general, to the field of computers and computer storage devices. More particularly, the present invention relates to efficient management of a redundant storage device array in a computing system.

Description of Prior Art

In computing systems designed for large data processing and data storage applications, redundant storage devices are provided to enhance the integrity of data maintained on the system in the event of a failure of a storage device.

For example, level 5 of RAID (Redundant Array of Inexpensive Disks; RAID-5) is a technology which utilizes an array of disk drives which contain data and parity information distributed, or striped, across each disk in the array. Parity information is additional, non-essential data used to reconstruct data contained on any of the drives of the array in the event of a single drive failure. In this manner, a RAID-5 disk array can improve the data integrity of the computing system by providing for data recovery despite the failure of a single disk drive.

A fundamental requirement of a RAID-5 system is that the data and the parity information must be synchronously maintained at all times in order to avoid data corruption. In other words, a RAID-5 disk array requires that the data and parity maintained within the storage system must be synchronized in order to correctly regenerate the data stored on any failed disk drive.

Maintaining synchronization between data and parity in a RAID-5 system becomes complicated when the RAID-5 system is implemented within a general purpose computing system, also known as a host-based computing system. In a host-based computing system, maintaining synchronization between the data and the parity information becomes more difficult because there are numerous events occurring therein which can potentially interrupt the synchronization between the storage of data and parity information.

In order to write new data to the RAID-5 array, conventional RAID-5 techniques involve numerous disk accesses to read, modify, and write the new data and the new parity to ensure that synchronization is maintained.

Because of the numerous disk accesses required, conventional RAID-5 storage techniques are generally characterized by slow processing times, when compared to non-redundant systems, for storing or writing data in the redundant array of disk drives. As mentioned, a single write of new data to the RAID-5 disk array generally requires six to eight disk input/output operations. Since a single disk input/output operation takes approximately ten milliseconds per operation, a single write of new data to a RAID-5 disk array can con-

ventionally require approximately sixty milliseconds.

Therefore, while a RAID-5 disk array improves the data integrity of the computing system in the event of a single disk failure, the performance of the RAID-5 disk array is very costly in terms of slower processing times.

SUMMARY OF THE INVENTION

In accordance with this invention, the above problems have been solved by a method for writing new data and new parity in a computing system having a system memory, an array of drives arranged in a RAID configuration, and a backup power module. A cache for storing parity and data information is established and maintained in the system memory of the host. By using the cache, the performance of the computing system in implementing a write of new data to the RAID disk array is substantially improved.

Specifically, the new parity and new data is calculated from the old parity and old data stored in the computing system. After the new parity is calculated, the new data is transferred to the cache for storage, and the new parity is transferred to the cache for storage. The computing system now has a stable version of the new data and the new parity.

The new data is then written to disk, and the new parity is written to the cache. Finally, the new parity transferred in the cache and the new data transferred to the cache are then marked as invalid.

System conditions of the computing system are monitored and if a fault is detected, the cache is disabled, and all data RAID-5 operations are performed using the disk only. The contents of the cache, along with a corresponding checksum, are copied to the drives so that the cache information can be used for data or parity reconstruction.

The above computer implemented steps in another implementation of the invention are provided as an article of manufacture, i.e., a computer storage medium containing a computer program of instructions for performing the above described steps.

In a machine implementation of the invention, an apparatus for storing data to a storage device configured as a RAID disk array in a computer, where the computer has a processor, an input/output device, a system memory, and a backup power supply for supplying power to the computing system. The apparatus comprises a memory cache, a checkpoint module, and a cache manager.

The memory cache, resident in the system memory, stores data. A checkpoint module detects a fault in the computing system and generating a fault indicator. The faults detected include a failure of software, a disturbance in the system's power supply, or a hardware fault.

A cache manager writes data and parity to the memory cache in a first mode, and writes data and parity to the storage device in a second mode. In response to the fault indication, the checkpoint module generating

a checksum of data contained in the cache, copying the data contained in the cache to the storage device, and switching the cache manager from the first mode to the second mode.

The great utility of the present invention is to substantially reduce the number of required accesses to a physical disk drive during a write operation to a RAID-5 disk array, thereby improving the performance of the computing system by reducing the time required to perform a single write operation.

Still another utility of the present invention is to maintain synchronization between the data and the parity information contained within the storage devices of the computing system.

Still another utility of the present invention is to provide synchronization between data and parity information in a host based implementation of a RAID-5 disk array.

Still another utility of the present invention is to provide recovery of data contained within a single disk drive which fails.

Still another utility of the present invention is to permit recovery of data or parity information after a single failure of other non-disk drive components of the computing system, such as a software failure or a CPU module failure.

The foregoing and other useful features and advantages of the invention will be apparent from the following more particular description of a preferred embodiment of the invention as illustrated in the accompanying drawings.

BRIEF DESCRIPTION OF DRAWINGS

FIG. 1 illustrates a computing system to perform the computer implemented steps in accordance with the invention.

FIG. 2 illustrates a block diagram of the preferred embodiment of the present invention.

FIG. 3 illustrates the logical operations performed to write new data in a RAID-5 disk array.

FIG. 4 illustrates the logical operations of to write new data to the RAID-5 disk array using a memory cache.

FIG. 5 illustrates an alternative embodiment of the cache director of FIG. 2.

FIGS. 6A and 6B illustrate the logical operations to determine if the memory cache should be enabled or disabled.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

The embodiments of the invention described herein are implemented as logical operations in a computing system. The logical operations of the present invention are implemented (1) as a sequence of computer implemented steps running on the computing system and (2) as interconnected machine modules within the comput-

ing system. The implementation is a matter of choice dependent on the performance requirements of the computing system implementing the invention. Accordingly, the logical operations making up the embodiments of the invention described herein are referred to variously as operations, steps, or modules.

The operating environment, in which the present invention is used, encompasses a stand alone computing system as well as the general distributed computing system. In the distributed computing system general purpose computers, workstations, or personal computers are connected via communication links of various types, in a client-server arrangement, wherein programs and data, many in the form of objects, are made available by various members of the system. Some of the elements of a stand alone computer or a general purpose workstation computer are shown in FIG. 1, wherein a processor 20 is shown, having an input/output (I/O) section 21, a central processing unit (CPU) 22 and a memory section 23. The I/O section 21 is connected to a keyboard 24, a display unit 25, a disk storage unit 26, network interface 30, and a CD-ROM drive unit 27. The CD-ROM unit 27 can read a CD-ROM medium 29 which typically contains programs 28 and data. The computer program products containing mechanisms to effectuate the apparatus, and methods of the present invention may reside in the memory section 23, or on a disk storage unit 26, or on the CD-ROM 29 of such a system. Examples of such systems include SPARC systems offered by Sun Microsystems, Inc., personal computers offered by IBM Corporation and by other manufacturers of IBM-compatible personal computers, and systems running the UNIX operating system or Solaris™ operating system.

Fig. 2 illustrates a block diagram of the preferred embodiment of the present invention. RAID-5 cache director 40 receives requests from the computing system to write new data to the array of disks 44. In Fig. 2, individual disks 45 are arranged in a RAID-5 disk array 44. From the perspective of the computing system, the disk array 44 appears as a single logical disk drive, although it is physically implemented as a plurality of disk drives in a RAID-5 array.

The operations of RAID-5 disk arrays are described in detail in the publication RAID: High Performance, Reliable Secondary Storage, by Peter Chen, et. al., published in ACM Computing Surveys, October 29, 1993.

RAID-5 cache director 40 implements operations to efficiently utilize the host memory cache 42 and the disk array 44. As will be described in detail below, RAID-5 cache director 40, responsive to a request to write new data, utilizes host cache memory 42 in order to reduce the number of operations involving the array of disk drives 44.

The RAID-5 cache director 40 also monitors system conditions in order to determine if cache 42 should be disabled. Backup power supply 46 provides uninterruptible power to the computing system including cache

director 40, cache 42, and the array of disk drives 44.

Host memory cache 42 is a cache which is created from the system memory of the host, typically in system RAM. Cache 42 can be permanently allocated from and reserved in the system memory of the host. By reserving a portion of system memory for cache 42, no other applications or devices can disturb the contents maintained in the memory space reserved for cache 42.

Cache 42 is stably maintained in traditional system RAM of the host through the use of a system backup power supply 46. In this manner, the contents of the cache could be copied to the array of disks, as will be explained below. The memory utilized for cache 42 is a matter of choice depending upon the system in which the preferred embodiment of the present invention is utilized.

The operation of RAID-5 cache director 40 in conjunction with cache 42, the array of disks 44, and backup power supply 46 are illustrated in Figs. 3, 4, 6A, and 6B.

Fig. 3 illustrates the logical operations performed by cache director 40 to write, without caching, new data to disk array 44. Operation 50 reads the old data from the disk, while operation 52 reads the old parity from the disk. Operations 50 and 52 are needed to calculate the new parity information. Operation 54 generates the new parity information by first removing the old data from the parity information. Removing the old data from the parity can be achieved through an exclusive-OR operation. The new parity information is then generated by including the new data into the parity information, which can also be achieved using an exclusive-OR calculation.

Having calculated the new parity information corresponding to the new data, it is advantageous, in order to maintain synchronization, to record the new data and the new parity in the disk before writing the new data and the new parity to its final location on the disk. In this manner, if the computing system is interrupted or if a device fails before the new data and new parity are written to the disk, the information is synchronized. As previously explained, synchronization between data and parity is needed to reconstruct data stored on a failed disk drive. Operation 56 records the new data to the disk, and then records the new parity information to the disk. This information is recorded to the disk in a known location so that if the system crashes before, during, or after operation 56, the computing system can access the known location of this information in order to reconstruct, if necessary, any data.

Having permanently recorded the new data and new parity, this information can now be transferred to their respective storage locations on the disk drives. Operation 58 writes the new data to the disk, and operation 60 writes the new parity information to the disk. In this manner, both the new data and the new parity are now synchronously maintained on the disk drive.

Having completed the write of the new data and the new parity to the disk, operation 62 marks the entire write operation as complete. By marking the operation

completed, the application requesting the write operation can then continue with its next desired operation. Furthermore, by marking the operation complete, the system is made aware that the data and the parity recorded by operation 56 can be discarded or overwritten. If a disk drive in the system should fail, the parity information written by operation 60 can be used to regenerate the data contained on the failed disk drive.

Fig. 4 illustrates the logical operations performed to write new data, using host memory cache 42 (Fig. 2), along with disk array 44. When the RAID-5 cache director utilizes memory cache 42, substantial reductions in processing time are realized because the memory cache has a substantially shorter response time than do the physical disks 45 of the disk array 44.

Responsive to a request to write new data to the disk array 44, operation 70 reads the old data from the disk. In order to read the old parity so that the new parity can be calculated, operation 72 determines if the old parity exists in the cache. If the old parity is not stored in the cache, operation 74 reads the old parity from the disk. Otherwise, operation 76 reads the old parity from the cache 42 (Fig. 2). As previously discussed, reading or writing data from the memory cache is substantially faster than reading or writing data from the physical disk drive.

After the old parity is fetched, operation 78 generates the new parity using the old data, old parity, and new data, as explained in operation 54 of Fig. 3.

Again, it is advantageous to store the new data and the new parity into a permanent location prior to writing this information to the final locations on the disk. RAID-5 cache director 40 (Fig. 2) again utilizes host memory cache 42 in order to enhance the speed performance of these operations.

Operation 80 transfers the new data to the memory cache. Operation 82 then transfers the new parity information to the memory cache. There are a variety of cache implementations available. For instance, a least recently used (LRU) link list can be used to organize the data and parity information stored in the cache. As will be explained below, the contents of the cache can be overwritten by a subsequent write operation only after the contents of the cache have been successfully transferred to disk.

Having stored the new data and new parity in the cache, operation 84 writes the new data to the disk at its appropriate storage location. Operation 86 then writes the new parity information to the cache. Operation 86 utilizes a different storage region in the cache than operation 82 for maintaining the new parity information. At the completion of operation 86, both the new data and the new parity have been stored consistently so that synchronization is maintained.

Operation 88 then invalidates the data and the parity stored in the cache corresponding to operations 80 and 82. In this manner, this information is marked as no longer being usable for data recovery purposes. In this manner, the memory occupied by the information stored

in the cache through operations 80 and 82 can be utilized by subsequent write operations executed in the computing system. It should be noted that the new parity information stored in the cache by operation 86 is unaffected by the invalidation operation 88, because operation 86 utilized a different area of the cache than operations 80 and 82.

Because the operations of Fig. 4 utilize the memory cache 42 (Fig. 2), the overall time required to complete operations 70-88 of Fig. 4 is substantially shorter than the time required to complete operations 50-62 of Fig. 3. As will be described below, the decision to utilize memory cache 42 is made by RAID-5 cache director 40 based on the system conditions (Fig. 2).

The cache (Figs. 2 and 5) can be periodically "pruned" so that the oldest parity stored in the cache by operation 86 (Fig. 4) is moved to persistent storage in the disk array 44. A separate process (not shown) could perform this function when appropriate.

Fig. 5 illustrates an alternative embodiment of the present invention. Checkpoint module 90, cache manager 92, and disk manager 94 are separate components which can comprise RAID-5 cache director 40. Checkpoint module 90 monitors system conditions and issues commands to cache manager 92. Cache manager 92, responsive to a request to write new data, utilizes host memory cache 42 and disk manager 94 to write data and parity either to the cache 42 or the array of disks 44, or both.

Backup power supply 46 provides uninterruptable power for the entire computing system including cache 42, disk array 44, checkpoint module 90, cache manager 92, and disk manager 94.

Fig. 6A illustrates the logical operations performed by checkpoint module 90. Checkpoint module 90 is designed to detect changes in system conditions which would require that the contents of host memory cache 42 be transferred to persistent storage in disk array 44. Furthermore, checkpoint module 90 determines whether caching of data and parity should be enabled or disabled based on system conditions.

Operation 100 determines if a software failure has occurred within the computing system. Operation 102 determines whether a disturbance in the power supply of the system has occurred. Operation 104 determines if any hardware faults have been detected. If any of the conditions tested in operation 100, 102, or 104 are true, then operation 106 generates a checksum of the contents of cache memory 42. Operation 108 then copies the contents of the cache, along with the checksum, to the disk array 44 (Fig. 5) for persistent storage. Operation 110 then disables caching in the RAID-5 cache director so that the memory cache, which may be volatile due to the conditions detected at operations 100, 102, or 104, is not used. Therefore, the RAID-5 cache director would satisfy write requests by utilizing the operations previously discussed and shown in Fig. 3.

As system conditions are monitored according to Fig. 6A, Fig. 6B illustrates the corresponding operations

implemented in the RAID-5 cache director 40. Operation 112 determines if caching has been enabled, and if so, operation 114 uses the memory cache 42 (Figs. 2 and 5) to satisfy write requests in the computing system. Operation 114 would invoke the operations 70-88 of Fig. 4, previously described, to reduce the time to complete a write of data to the RAID-5 disk array 44 (Figs. 2 and 5).

If caching is disabled as detected at operation 112, then operation 116 satisfies the write operation by invoking the operations 50-62 of Fig. 3, as previously described. The write operations would be performed using the RAID-5 disk array 44 (Figs. 2 and 5) exclusively.

In the event of a single disk failure in the computing system, the data stored on the failed disk can be reconstructed because the operations of Fig. 3 and Fig. 4 have maintained consistency of and synchronization between the data and the parity.

Furthermore, data or parity information can be recovered after a single failure in a non-disk drive component of the computing system (i.e., software error, CPU module failure). Assuming these failures require that the system be rebooted and the failed component is replaced, upon reboot the disk drive which would contain the cache, in accordance with operation 108 of Fig. 6A, is read to determine if the cache was saved by the checkpoint module 90 (Fig. 5, 6A). If the cache was saved to the disk drive, then the contents of the cache can be used for data recovery as previously described. If the cache was not saved, as in the case of a major hardware error such as a CPU module failure, then the parity information can be regenerated from the data stored on the array of drives 44 (Figs. 2, 5). While regeneration of parity from the data stored in the drive array is a lengthy process, the parity information in the computing system can be restored.

Therefore, the present invention provides for a host based implementation of RAID-5, using a memory cache maintained in the host memory, to substantially reduce the amount of time required by the computing system to satisfy a request to write new data. The integrity of the RAID-5 information maintained in the system is achieved through the use of an uninterruptable power supply in conjunction with a check point module which monitors system conditions.

While the invention has been particularly shown and described with reference to preferred embodiments thereof, it will be understood by those skilled in the art that various other changes in the form and details made by made therein without departing from the spirit and scope of the invention.

Claims

1. In a computer having a processor, an input/output device, a system memory, a backup power supply (46) for supplying power to the computing system, and a storage device configured as a RAID disk

array (44), an apparatus (40) for storing data to the storage device, said apparatus comprising:

a memory cache (42) storing data, said memory cache resident in the system memory of the computing system; 5
a checkpoint module (90) detecting a fault in the computing system and generating a fault indication;
a cache manager (92) writing data and parity to said memory cache in a first mode, and writing data and parity to said storage device in a second mode; 10
in response to the fault indication, said checkpoint module (90) generating a checksum of data contained in said cache (42), copying the data contained in the cache to said storage device (44), and switching said cache manager from the first mode to the second mode. 15

2. The apparatus of claim 1, wherein the fault detected by said checkpoint module (90) includes a failure of software (100), a disturbance of power (102), or a hardware fault (104). 20

3. A method for writing new data and new parity in a computing system having a system memory, an array of drives arranged in a RAID configuration (44), and a backup power module (46) for supplying power to the computing system, the method comprising the computer-implemented steps of: 25

providing a cache (42) in the system memory for storing parity and data information;
calculating the new parity (78) and new data from old parity (74, 76) and old data stored in the computing system; 30
transferring (80) the new data to the cache for storage therein;
transferring (82) the new parity to the cache for storage therein; 40
writing (84) the new data to the disk;
writing (86) the new parity to the cache; and
invalidating (88) the new parity and new data stored in the cache. 45

4. The method of claim 3, further comprising:

detecting (100, 102, 104) if the cache should be disabled based on an operating condition of the computing system. 50

5. The method of claim 4, wherein the detecting step further comprises the steps of:

generating (106) a checksum of the contents of the cache; 55
storing (108) the contents of the cache to at least one drive of said array of drives; and

disabling (110) use of the cache for data storage responsive to said detecting step.

6. A computer program storage medium readable by a computing system and encoding a computer program of instructions for executing a computer process for writing new data and new parity in a computing system having a system memory, an array of drives arranged in a RAID configuration, and a backup power module for supplying power to the computing system, said computer process comprising the steps of:

providing a cache in the system memory for storing parity and data information;
calculating the new parity and new data from old parity and old data stored in the computing system;
transferring the new data to the cache for storage therein;
transferring the new parity to the cache for storage therein;
writing the new data to the disk;
writing the new parity to the cache; and
invalidating the new parity and new data stored in the cache.

7. The computer program storage medium of claim 6 where the computer process further comprises the steps of:

generating a checksum of the contents of the cache;
storing the contents of the cache to at least one drive of said array of drives; and
disabling use of the cache for data storage responsive to said detecting step.




Fig. 1


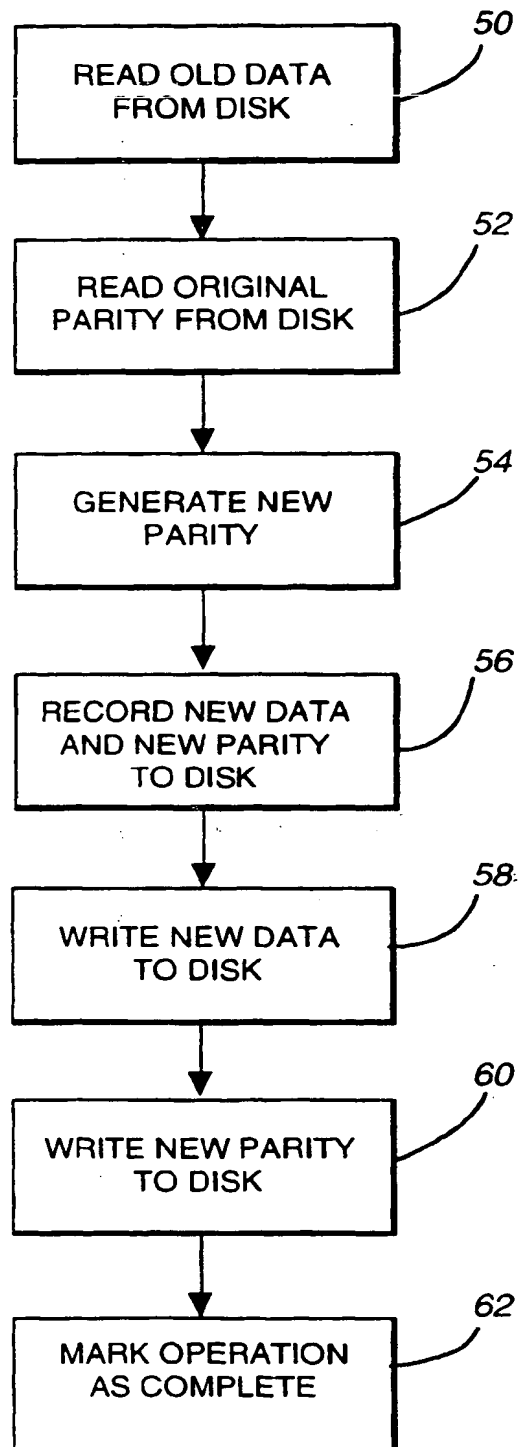
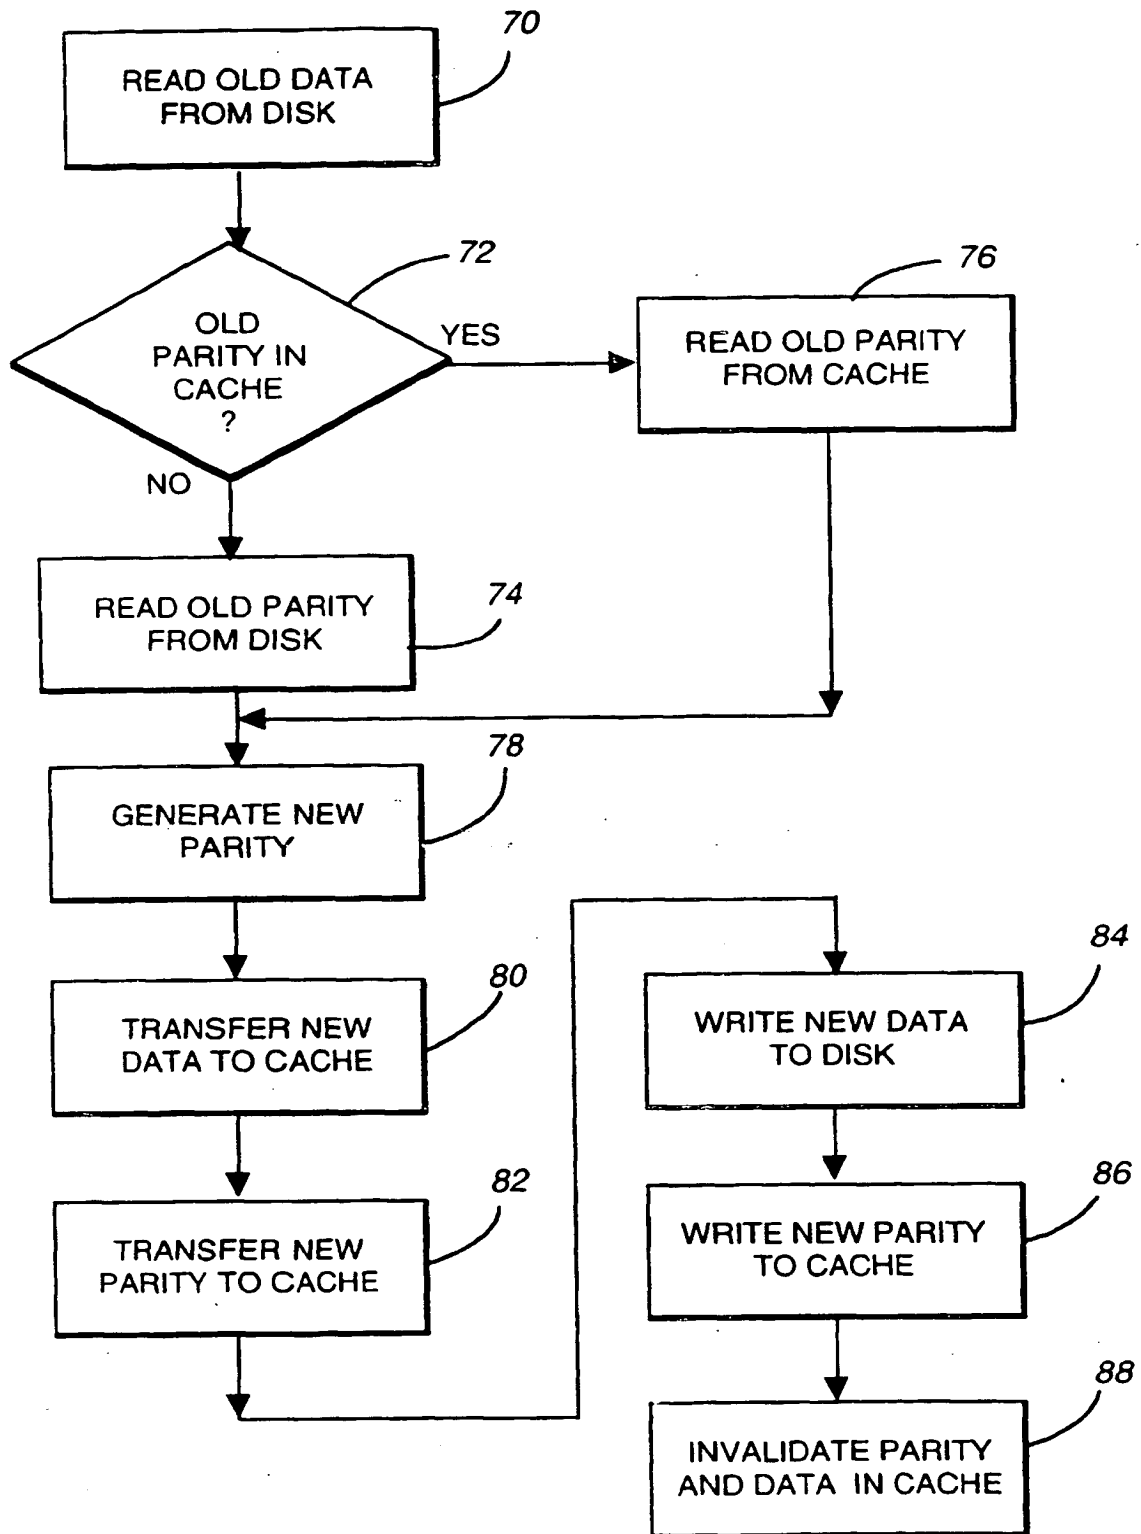


Fig. 2

**Fig. 3**

**Fig. 4**


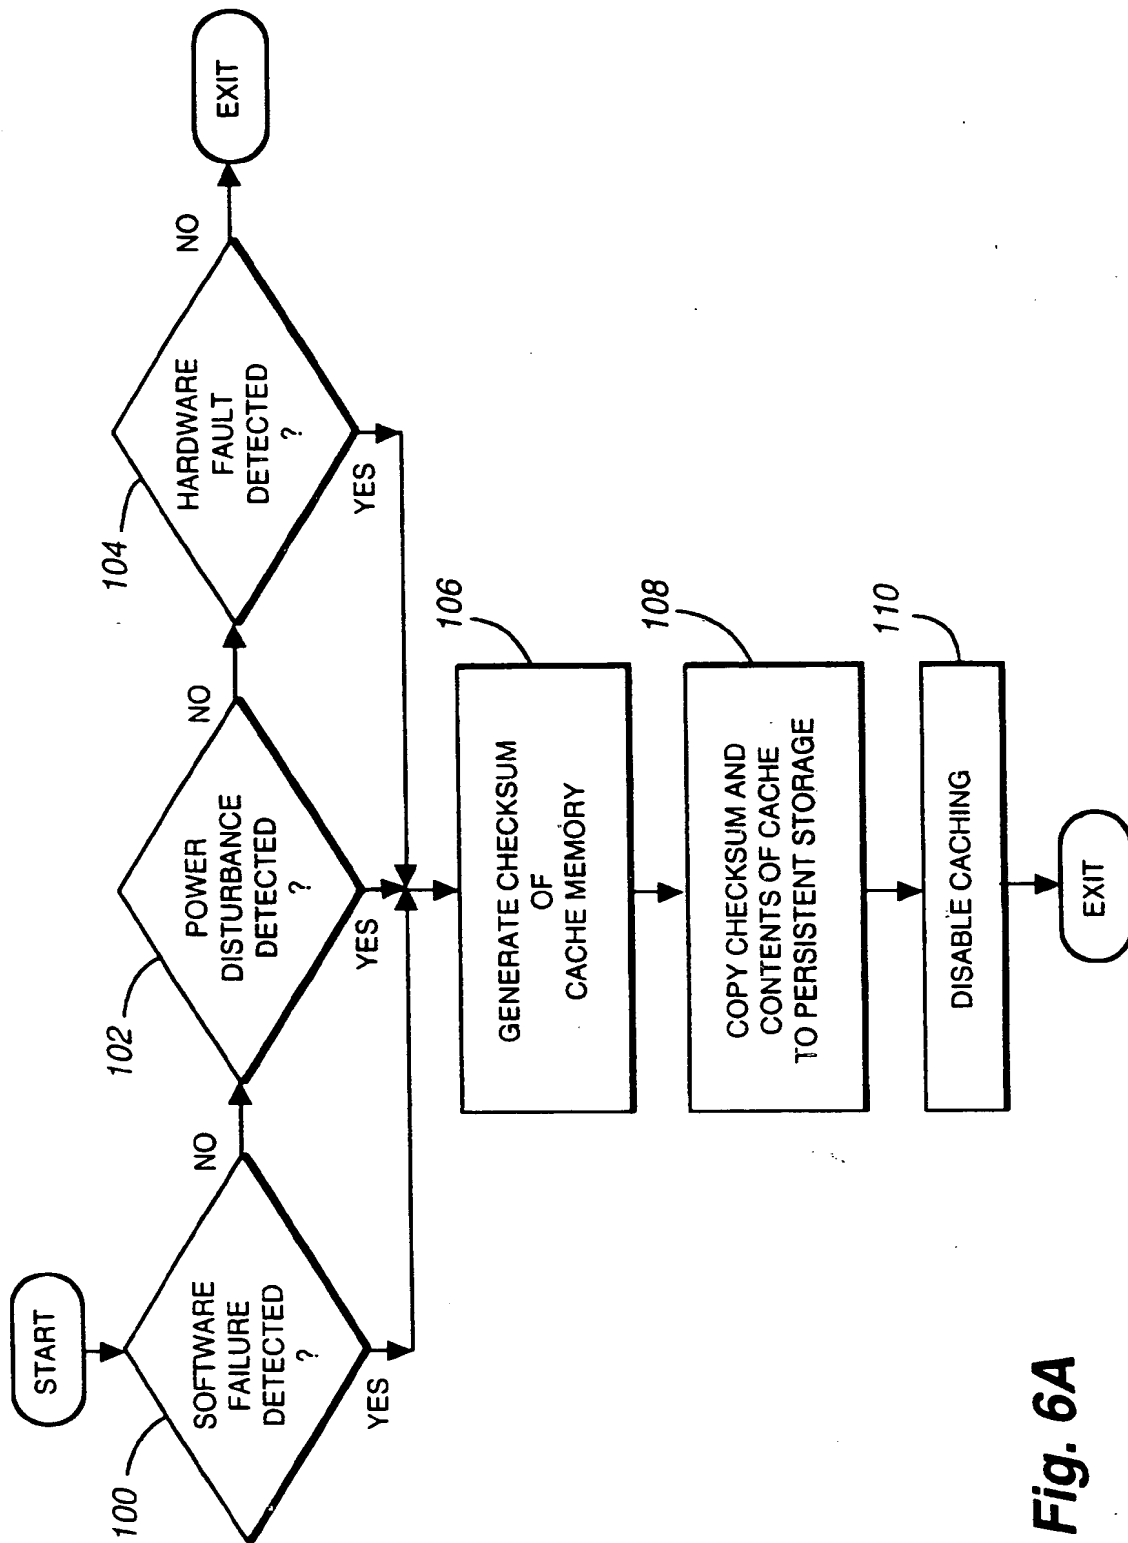
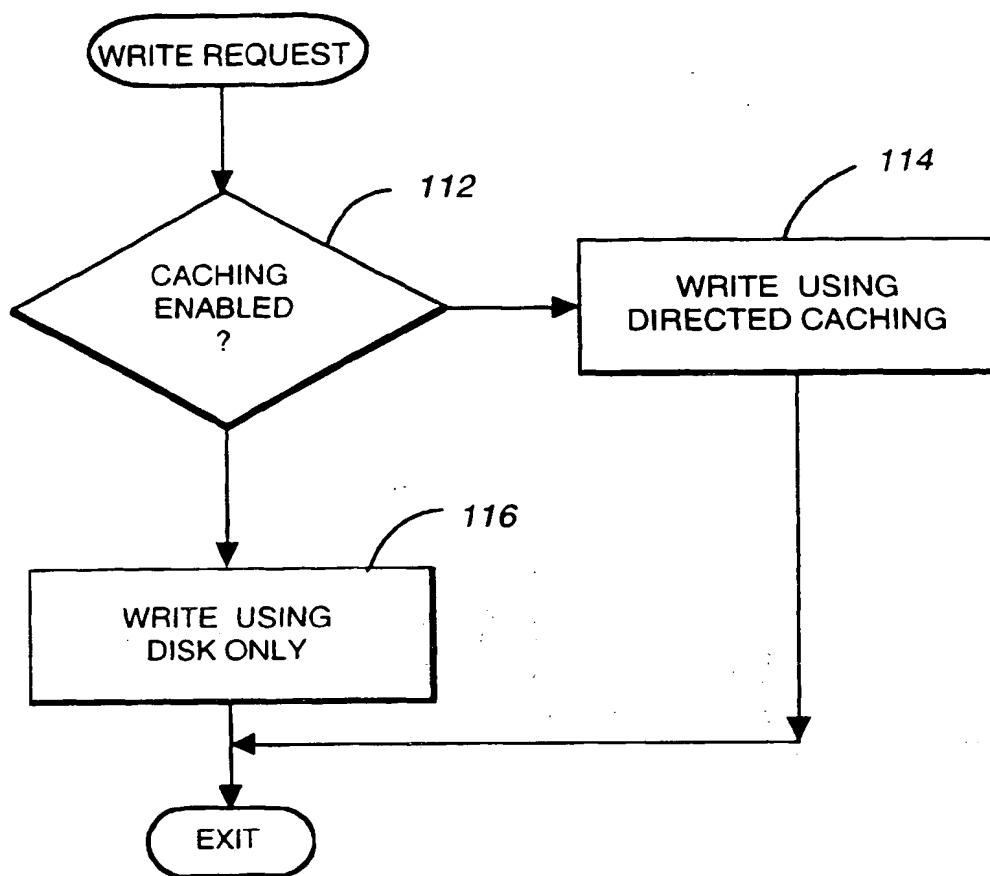


Fig. 5

**Fig. 6A**

**Fig. 6B**



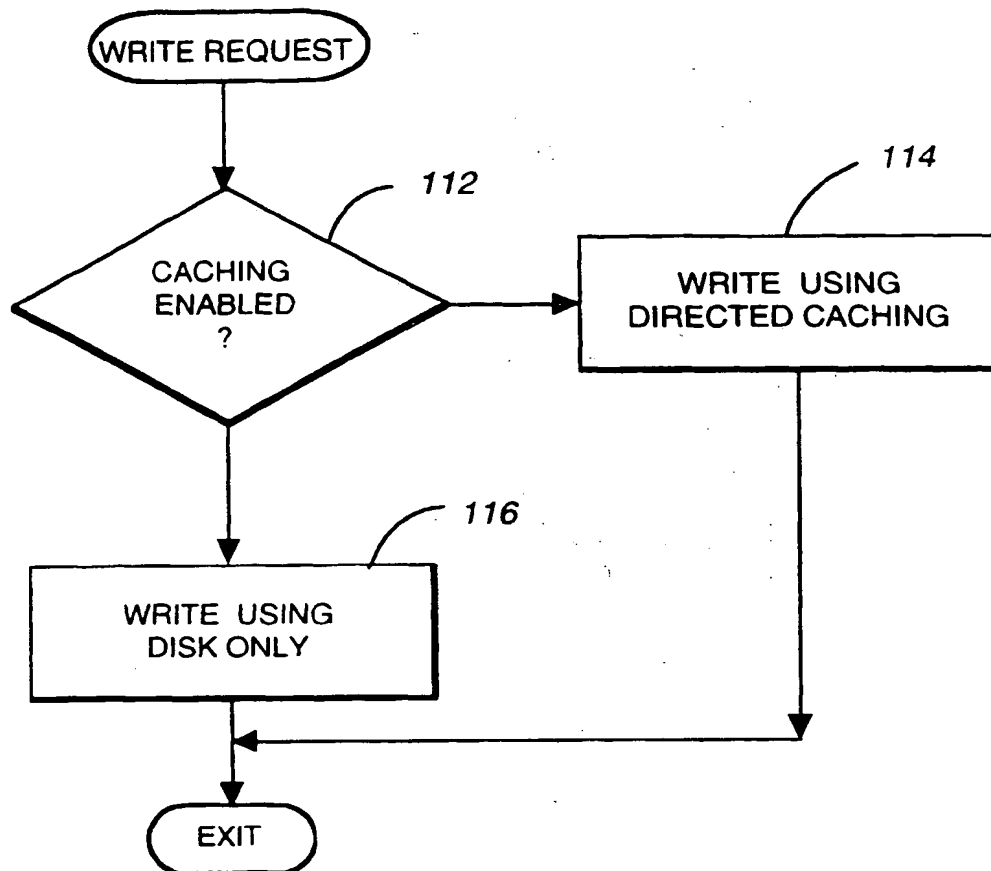
European Patent
Office

EUROPEAN SEARCH REPORT

Application Number
EP 97 10 4950

DOCUMENTS CONSIDERED TO BE RELEVANT			
Category	Citation of document with indication, where appropriate, of relevant passages	Relevant to claim	CLASSIFICATION OF THE APPLICATION (Int.Cl.6)
A	EP 0 569 755 A (MITSUBISHI ELECTRIC CORP) 18 November 1993 * abstract * * column 2, line 2 - line 15 * * column 4, line 8 - column 5, line 5 * * column 5, line 35 - line 50 * * figures 4,5 *	1,3,6	G06F12/08 G06F11/10
A	--- COMPUTER ARCHITECTURE NEWS, vol. 21, no. 2, 1 May 1993, pages 76-86, XP000380356 JAI MENON ET AL: "THE ARCHITECTURE OF A FAULT-TOLERANT CACHED RAID CONTROLLER" * the whole document *	1,3,6	
A	--- COMPUTER TECHNOLOGY REVIEW, vol. 15, no. 6, June 1995, page 44, 46/47 XP000516565 MASSIGLIA P: "WRITE-BACK CACHE AND RAID: CACHE AND I/O PERFORMANCE" * the whole document *	1,3,6	
			TECHNICAL FIELDS SEARCHED (Int.Cl.6)
			G06F
The present search report has been drawn up for all claims			
Place of search BERLIN		Date of completion of the search 20 June 1997	Examiner Masche, C
<p>CATEGORY OF CITED DOCUMENTS</p> <p>X : particularly relevant if taken alone Y : particularly relevant if combined with another document of the same category A : technological background O : non-written disclosure P : intermediate document</p> <p>T : theory or principle underlying the invention E : earlier patent document, but published on, or after the filing date D : document cited in the application L : document cited for other reasons A : member of the same patent family, corresponding document</p>			

EPO FORM 150 (12/92) (P04C01)

**Fig. 6B**

THIS PAGE BLANK (USPTO)